From Science to Solutions

Cloud Computing Cluster

Introduction to Cloud Computing

Rick Martin, Co-chair, Cloud Computing Cluster
August 26, 2013

Senior IT Strategist
SAIC

What is Cloud Computing?

- Cloud computing is the convenient, rapid access to a shared pool of computing resources.
- Such resources can be strictly infrastructure (networks, servers, storage, etc.) or can be combined with software to facilitate ready access to applications and services.
- Cloud resources offer attractive flexibility—as demand ebbs and flows the amount of 'horsepower' being consumed can be adjusted up or down.
- Cloud computing is not resource virtualization with a new name; features such as self-service provisioning and advanced metering of use set it apart as a new and transformational technology.

NIST Cloud Characteristics

On-demand self-service
- A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.

Broad network access
- Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling
- The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).
- Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid elasticity
- Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in.
- To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time (incrementally as required).

Measured Service
- Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

Deployment Models

- **'Software as a service'** (SaaS) is where the provider offers an application for use by the customer. Web-based email from Google being deployed for NOAA is an excellent example.
- '**Platform as a service'** (PaaS) is when a customer application runs on a software platform or stack maintained by the provider. For example, a customer owned application could be run in the cloud utilizing a ArcGIS or Oracle database instance offered by the provider.
- **'Infrastructure as a service**' serves those who simply wish to acquire computing horsepower on demand and do not wish to maintain it locally. Such capability tends to be viewed as an extension of the local data center and engenders more direct involvement from IT staff.

Deployment Types

- A **private cloud** is devoted to a single customer and the resources are managed for their sole benefit. Private clouds can be a cluster of networked servers directly owned by the organization or a dedicated, segmented, and protected area on a provider's network.
- A **public cloud** provider offers service to a range of customers across government and industry boundaries without regard to the nature of the data involved. Among cloud solutions, this is the lowest cost option as the provider maximizes their economies of scale and passes that along to their customers.
- A **community cloud** offers benefits of both the private and public clouds. Essentially, it is a public cloud whose tenants are limited to a defined group or class of customers. By limiting the tenants, the provider may tailor the offering to specific needs and demands of that group.
- A **hybrid cloud** is a mixture of one or more of the above.

Taxonomy of Clouds

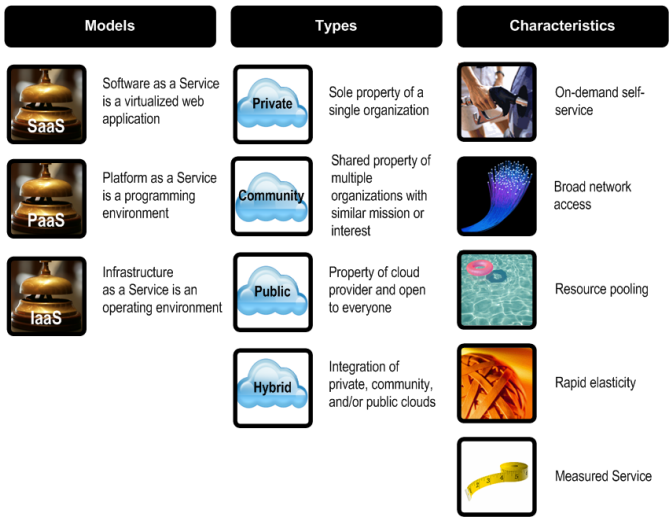

Scenario 1: Simple Private Cloud

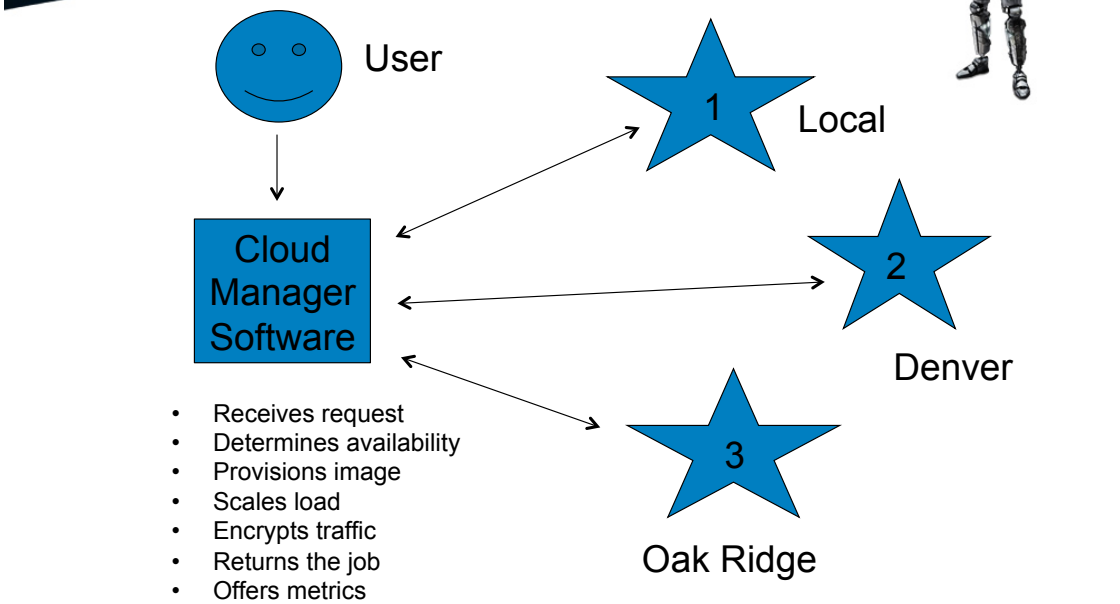

- Receives request
- Determines availability
- Provisions image
- Scales load
- Encrypts traffic
- Returns the job
- Offers metrics
- Allocates the costs

Scenario 2: Public or Community/ IaaS

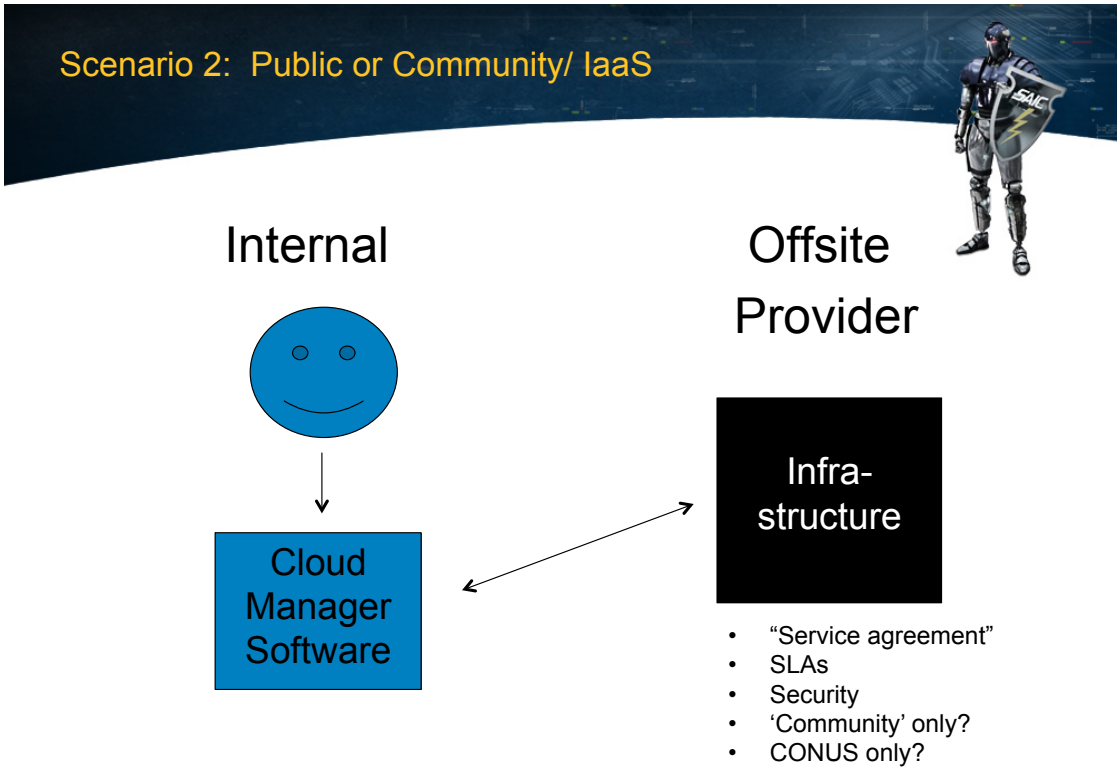

- "Service agreement"
- SLAs
- Security
- 'Community' only?
- CONUS only?

Scenario 3: Public or Community/ PaaS

Internal

Offsite Provider

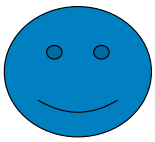

Cloud Manager

Platform

Infrastructure

Scenario 4: Public or Community/ SaaS

Internal

Offsite Provider

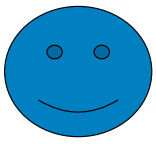

Website

Scenario 5: Hybrid

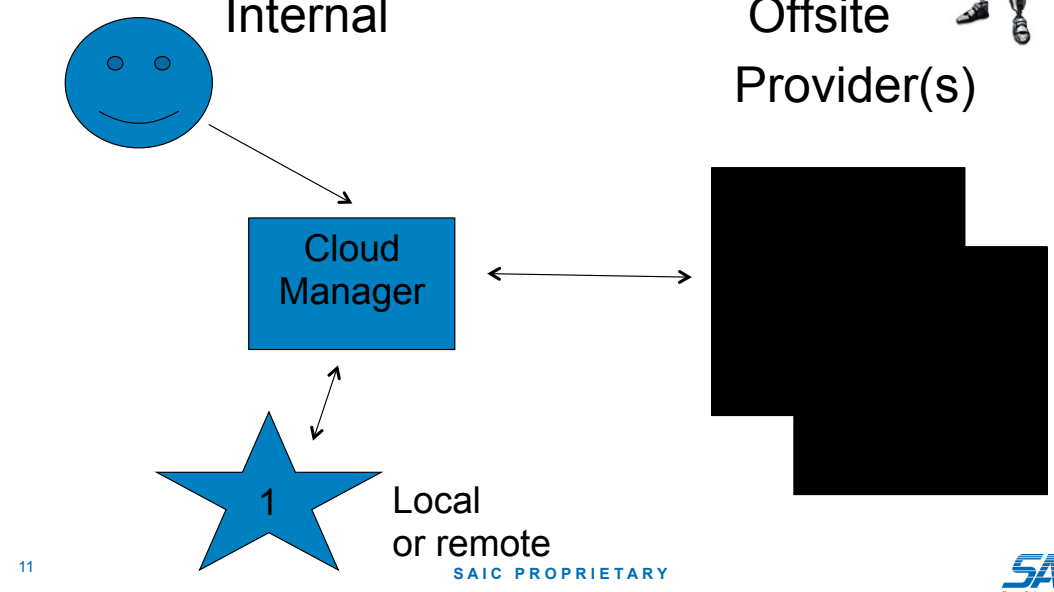

Storage, Backup and Recovery

- Storage costs decrease; storage demand increases
- Internal storage cost comparisons must capture:
  - technology assessment, growth forecasting, acquisition and sys admin functions
  - telecom, , capital replacement, and backup/recovery.
- External storage costs are usually on a ‘schedule’ and include basic backup and recovery within the providers infrastructure
- Axiom is that storage should be closest to users; telecom can add latency and be a big cost factor.
- Metadata should be kept separate from data
- Archiving and backup are not substitutes for records management.

Cost Pools for Cloud Infrastructure

Cloud Brokerage

What is a cloud broker?

- Entity that manages the use, performance and delivery of cloud services
- Negotiates relationships between cloud providers and cloud consumers.
- Can be passive (merely an advisor) or active (directly playing a role)

Why cloud brokerage?

- Get the best out of the marketplace
- Manages multiple cloud and virtual resources
- Establish common management tools and facilitate cross-cloud services

Features of Cloud Broker

- Provides advice on selecting and negotiating with cloud provider
- Offers turn-key solution(s) which makes cloud adoption and use easy for customer
- Offers some standard solution templates and suggestions as to which vendors might be best
- Delivers pre-architected solutions from own catalog (augmented by other vendors)
- Maintains a inventory of vendors and services; recommends providers and suggests bundling to customer when requested
- Delivers full-featured, value-added cloud solutions itself and has an ecosystem of additional vendors with pre-negotiated service/prices.

Cloud Security and Reliability

- All basics of IT security apply to hosting data in a cloud
  - Access controls
  - Encryption
  - Backup and recovery
  - Reporting
  - Auditing
- GSA contract offerings have or will have undergone C&A
- All privacy and recordkeeping requirements apply as well
- 'Key' management an emerging issue; persistent PKI solutions under development
- Reliability questioned by many when cloud outages make news.

FedRAMP In One Slide

Promoting a Standardized Cloud Security Baseline

- Aims to establish a "do once, use many times" framework that saves cost, time and staff to conduct redundant security assessments
- Benefits include increasing re-use of existing security assessments across agencies, improving real-time security visibility, and supporting risk-based security management.

Program Phases

Launch

Pre-Launch
- Policy Memo
- 3PAO Applications
- Publish Requirements and Compliance Guidance

Initial Operating Capabilities
- 3PAO accreditation continues on a rolling basis
- First CSP provisional authorization from JAB issued December 26,2012

Full Operations

As of August 1, 2013:
- 22 3PAOs have achieved accreditation.
- ATOs have been issued for 8 cloud services.

Sustaining Operations
- *Initial program schedule anticipated scaling to steady operations in FY 2014.*

Agencies with existing CSPs/ATO have 2.5 years from June 2012 launch to migrate to FedRAMP compliant solutions.

According to the initial FedRAMP phase schedule, provisional authorization was expected to be issued to fewer than 20 vendors before entering Full Operation during Q2 of FY 2013.

Cloud Computing for Earth Science

- **'Big Data'**
- **Database Architectures and Trends**
- **Planning Considerations for Data in the Cloud**

Earth Science Datasets are 'Big Data'

"the tools, processes and procedures allowing an organization to create, manipulate, and manage very large data sets and storage facilities." ZDnet

- Sensor data
- Web and other logs
- Temporal and spatial data sets
- High-def video and image libraries

'Big Data' Variables

- Scale of Preserved Data
  - Original data
  - Derived and generated data
  - Need file system that can handle (Hadoop, etc)
- Variety
  - Geospatial data
  - Transaction data
  - Social media feeds
- Velocity
  - Automated telemetry
  - Collected 24/7
  - Distributed users

‘Big Data’ Conceptual Architecture: SAIC View

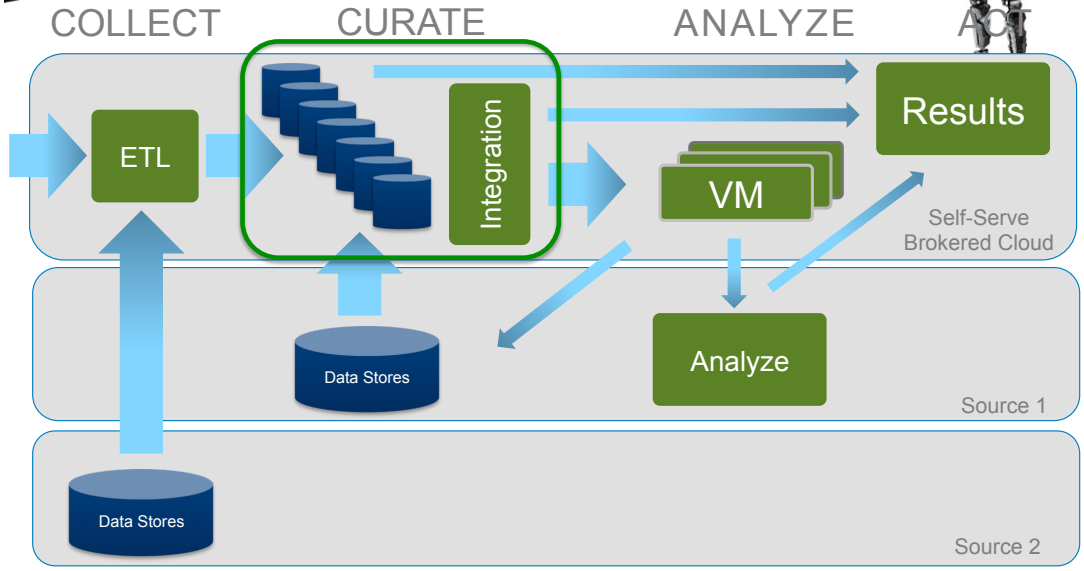

New Database Realities

- Most relational dbs are unsuited for analysis
  - Challenged by multi-tenancy
  - Less valuable for spatial and temporal analysis
  - Too costly to operate when large user-base runs complex queries
- Cost projections must assume long archival periods
- Data must be immutable with ironclad bind to metadata
- Focus on simpler and more scalable data structures that are effectively ‘giant tables’
- Abandon or minimize SQL (NoSQL)
- Cloud providers integrate these into their service offerings--PaaS
  - Google AppEngine data store
  - Amazon SimpleDb

Scalable Key Value File Systems: Hadoop

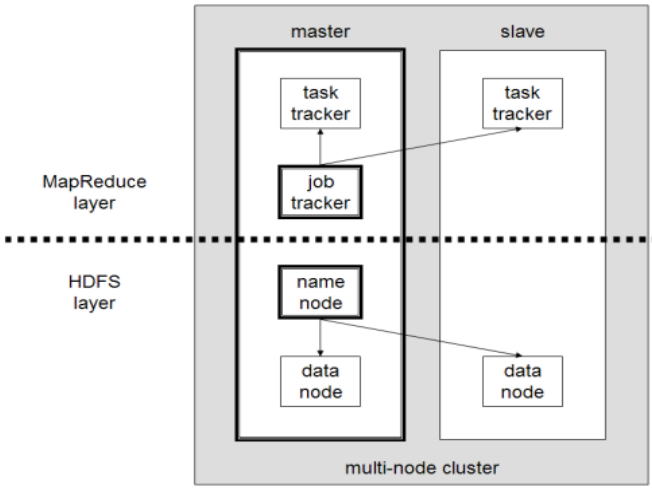

Scaling for Demand: New 'Slave' ~~Instance Created~~

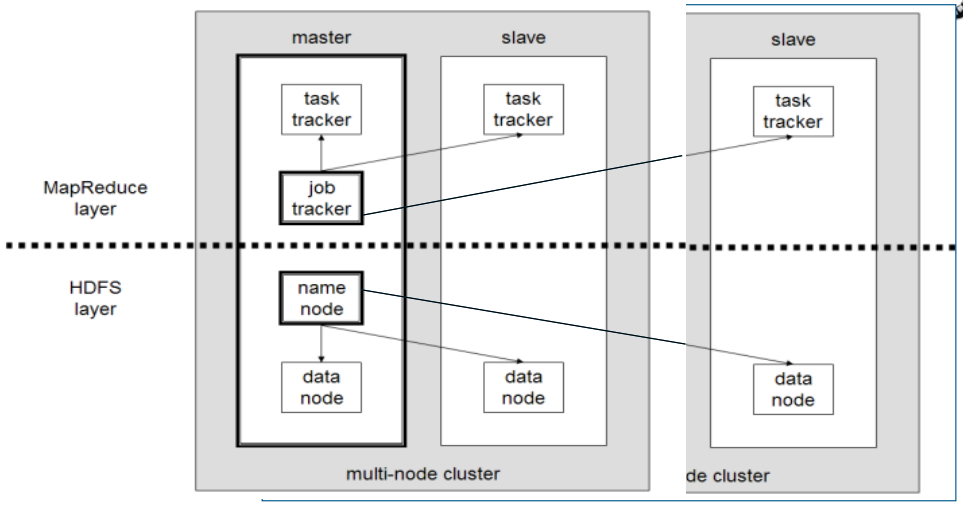

Planning for Data In The Cloud

What are the goals for implementation?
Which of the process, investment, value and intangible benefits of CC does the organization want to achieve?
What is the long-term vision for the preservation of data? Is there a retention schedule?
What business needs are driving the placement of data in a cloud and what is the timeline?
Is there an internal charge-back mechanism that needs to be 'fed'?

Which security, privacy, and continuity considerations will affect the design of your implementation?
How can security facilitate access and ensure preservation?
What level of availability is desired? Is that data publically available?
Does the security plan apply to the original data only? Or will there be derived data as well?

Do existing infrastructure and vendor relationships impact the design of a solution?
Is there an installed base of in-house technology that could benefit by being included in the solution?
Does data require an applications which uses a specified computing platform, both hardware and software?
Is there another reason—like the knowledge base of the users which suggests the inclusion of a particular technology?

What level of cloud management does the organization wish to directly undertake?
Who will be ordering and monitoring service—end users or the CIO staff on their behalf?
What skill mix does the CIO staff bring to the table or are willing to develop? The end users?
Does the CIO wish to offload infrastructure management to an external party?

Mobility is a Key Consideration in Cloud Architectures

To meet the 24x7 pressure

- Devices
  - Differing form factors
  - Differing interactions
  - Collection and retrieval
- Anywhere access
- Always connected
- Syncing









Spatial Cloud Computing

- *The computing paradigm that is driven by geospatial sciences, and optimized by spatiotemporal principles*
- *Enables geospatial and other science discoveries within distributed computing environment.*

Yang C., Goodchild M., Huang Q., Nebert D., Raskin R., Xu Y., Bambacus M., Fay D., 2011 (in press), Spatial Cloud Computing: How geospatial sciences could use and help to shape cloud computing? International Journal on Digital Earth

- Shared, cloud-based infrastructure composed of key geospatial services and applications.
- Geospatial services that can be built once and used many times
- May consist of infrastructure, data, metadata, and applications/tools based on commonly accepted standards
- Most geoplatforms will including an inventory of existing and planned resources will be utilized as a starting point for identifying key resources that may be included as the Platform evolves.
- Examples:
  - GIScloud.com
  - Google Cloud Platform
  - OpenGeo.org
  - ArcGIS Platform
  - MobiCloud

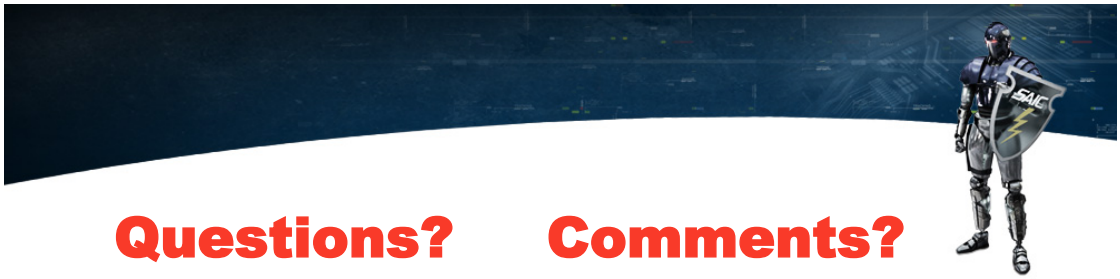

Questions? Comments?

For more information:

Rick Martin, Senior IT Strategist
richard.a.martin-2@saic.com
919-714-2527